CSinParallel Pacific Northwest 2016 Workshop

Introduction to CSinParallel

Dick Brown, St. Olaf College

Macalester College
Calvin
Minds in the Making
St. Olaf College

Workshop site

On CSinParallel.org:

http://serc.carleton.edu/csinparallel/workshops/PacificNW16/index.html

See also workshop handouts

Take-home messages

- The **traditional CS curriculum is "striking out"** against the fastballs and curve balls of current computing technologies
- Our students need **preparation in parallel and distributed computing (PDC)** for their careers
- The **CSinParallel strategy** for effective, quicker *curricular change*
- **Strategic resources** for teaching PDC

The BIG challenge: **Forming an effective community to support this curricular transition**

Striking out against new PDC technologies

http://armchairgm.wikia.com/File:1206127655_Strikeout.gif

Striking out against new PDC technologies

Strike 1: **Multicore processors**

- The norm for about **10 years**
- *Why?* "Hitting the wall" in:
  a. (Electrical) **power**
  b. **ILP** (can't hide much more parallelism within a core)
  c. Deepening **memory hierarchy**

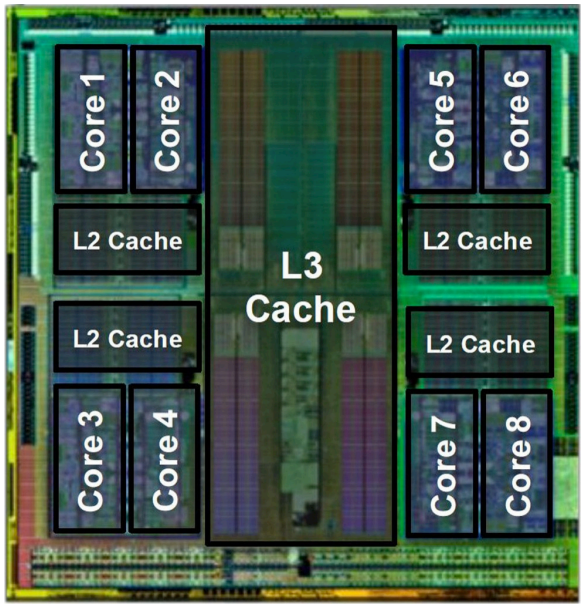

http://images.bit-tech.net/content_images/2012/11/amd-fx-8350-review/piledriver-3b.jpg

Striking out against new PDC technologies

Strike 2: **Cloud computing revolution**

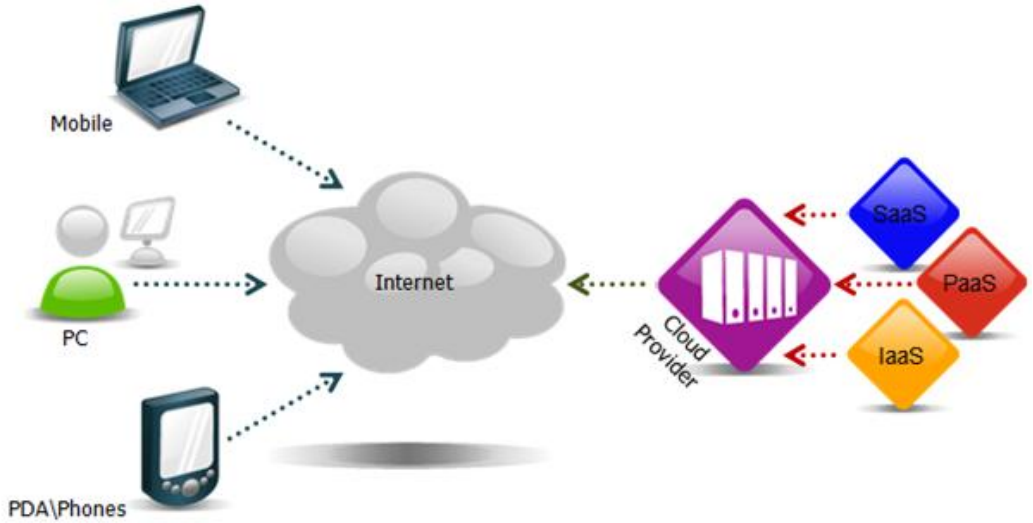

- Distributed computing empowers exciting new web services

http://cloudcomputingadvices.com/wp-content/uploads/2012/08/cloud_computing-Features.jpg

Striking out against new PDC technologies

Strike 3: **Heterogeneous computation**

(It's not just for HPC anymore)

- Commodity chips with multiple core types
  - Intel vector cores; AMD GPU + CPU cores
- Multicore and heterogeneous multicore as distributed/cluster computing nodes

Educating students for their careers

- Parallelism and concurrency traditionally taught in Architecture and Operating Systems
  - **Still necessary, just no longer sufficient**
- **Commodity computing**, not just HPC research
- **Hardware *and* software** evolving rapidly
  - Including languages, libraries, frameworks
- Recent **curriculum recommendations**
  - ACM/IEEE CS2013, for undergraduate CS majors
  - TCPP recommendations for PDC

Educating students for their careers

- Parallelism and concurrency traditionally taught in Architecture and Operating Systems
  - **Still necessary, just no longer sufficient**
- **Commodity computing**, not just HPC research
- **Hardware *and* software** evolving rapidly
  - Including languages, libraries, frameworks
- Recent **curriculum recommendations**
  - ACM/IEEE CS2013, for undergraduate CS majors
  - TCPP recommendations for PDC

… Rapid curricular change??

The CSinParallel strategy

1. Brief, flexible PDC teaching modules

- Almost any CS course, at almost any level
- 1- to 3-day course units, for feasible incremental modifications to a syllabus
- Broad variety of topics, technologies, languages, etc.
- Adaptable/editable for local modification
- Emphasis on hands-on exercises with current technologies
- Learning objectives, teaching tips, etc.

The CSinParallel strategy

2. Pedagogical effectiveness

- ***Small interventions, big impact***
  - Opportunities for broad exposure to PDC
- Effective ***hands-on learning*** with PDC tools
- ***Spiral approach***
  - Recurring topics in multiple contexts leads to better retention and deeper understanding
- Early and often
  - Message: ***PDC is natural and pervasive in CS***

The CSinParallel strategy

3. Community of folks seeking to teach PDC

- ***Community support***
  - Communicating with others trying same things (modules, institution types, courses, techs,…)
  - Resources for recording others' past experiences (e.g., Piazza, teaching tips)

Creating community/human networking is a primary goal for this workshop event

Some strategic resources

- **Modules**
  - **Basic CSinParallel modules**
    - Examples: WMR for CS1; Multicore Programming (intro); Concurrent Data Structures (C++ or Java); Parallel Sorting
  - **"Exemplar" modules**
    - Present a significant domain application (e.g., drug design, epidemiology, traffic flow) together with a sequential implementation
    - Choice of parallel/distributed implementations
    - Pedagogical possibilities **+** student motivation
  - **"Taste" modules**
    - Brief introductory experience of a new technology, or parallel language, larger educational work, etc.

Some strategic resources

- Relating local courses and CSinParallel materials to **latest curriculum recommendations** (CS2013, TCPP)

- **Parallel Programming Patterns**
  - Recurring design strategies for parallel programs, distilled from practices of experienced pros
  - Guides to problem solving and parallel thinking for undergraduate learners of PDC

Some strategic resources

- **Platform resources**
  - Example: *WebMapReduce (WMR)*, for beginning or advanced students to learn about scalable computations that fuel cloud-powered services
  - Example: *CDER hardware resources*, free educational access to PDC hardware
  - Example: *Inexpensive micro-clusters*
- **Students helping** profs, other students
  - Scalable collaborative resource
  - Student capacity for **exploring new technologies**

This workshop

Adaptable program – let us know *your* interests

*Basic plan focuses on **hands-on experience***

Tuesday	1. CSinParallel (✓); PDC; resources ***2. OpenMP*** (shared-mem. parallel) ***3. WebMapReduce*** (XXXXL data)
Wednesday	***1. MPI*** (distributed computing) 2. Some current trends ***3. GPU*** (SIMD parallel)
Thursday	1. Integrating PDC in *your* curriculum 2. Next steps?

Refined BIG challenge

How can we form a supportive community that **benefits professors more than it "costs"**

What do we mean by "costs"?

- **Learning new systems** (e.g., MPI, piazza)
- **Competition for time** during the term
- Weight of **inertia against change**

Recap

- Hardware and software parallel and distributed computing (PDC) technologies are **evolving rapidly, with no end in sight**
  - Multicore; cloud/distributed; *heterogenous*
- **Students need to know about PDC** as they enter the workforce.
- *CSinParallel offers modules, resources, and strategies for teaching PDC, in a context of* ***supportive community***
- Biggest unsolved problem: ***Creating a support community worth the effort***

Questions?